

Description

Method for redundant data management in computer networks

5

Technical field

The invention relates to a method for redundant data management in a computer network, in which data is consistently stored in at least two computers of a network.

10

Prior art

Computer networks have developed over the last few years to become a significant means of communication, by means of which 15 a plurality of services is offered. The most well-known computer network in the world is the internet, which denotes the totality of all networks linked together and the networks using the Internet Protocol IP as a transport protocol.

20 On level 3 of the OSI layer model, the Internet Protocol IP is responsible for the connectionless transport of data from a sender, if necessary over a number of networks, to the recipient, with neither error identification nor error correction taking place.

25

Transport protocols such as the Transmission Control Protocol TCP or User Datagram Protocol UDP are based on the Internet Protocol.

30 The Transmission Control Protocol TCP is a connection-oriented transport protocol, which enables a logical, reliable full-duplex-point-to-point link. This ensures that data packets are transmitted in an error free manner and in the desired

sequence over an underlying network according to the Internet Protocol IP.

The User Datagram Protocol UDP enables a connectionless point-to-point communication. In this way, the individual data packets (datagrams) are independent of one another.

Transmission is not protected. This means that data packets can get lost and that the receive sequence does not have to correspond to the send sequence. Any protection for data transport has, if necessary, to be provided at application level. UDP is generally used for simple services, in which a simple response (a datagram) to a simple request (a datagram) is expected.

One of the most popular applications of the internet is the World Wide Web WWW, which enables easy access to multimedia data on remote WWW servers by way of hyperlinks. The data which could comprise texts, audio sequences, fixed-images and films is viewed and displayed by means of WWW browsers, or local front-end software. The World Wide Web is based on the client/server principle, on which all important applications in the internet, FTP or News as well as WWW, are based.

Applications designed according to the client/server principle use a client program on the user side, said client program exchanging data with a specific service computer in the network, known as the server. In this way, the server is generally responsible for data management, while the client deals with the presentation of this data and interaction with the user. Client and server use a precisely defined protocol for this purpose.

Internet servers are operated either by commercial or non-commercial organizations, such as for instance universities,

which offer different online services, such as information services (weather reports, train timetables, message services) and communication services (e-mail, chat, forums), which are realized by means of corresponding application programs.

5

One essential criterion for the quality of a service offered is the probability, with which a potential user can use the service, the corresponding measured value, resulting from the downtime during a specific observation time, the so-called 10 availability of this service, having a significant influence on the satisfaction of the user and thus on the success of the respective service.

To increase availability, it is often requested that services

15 and/or the respective underlying data should be protected against failures of the relevant hardware components of the computer network, in other words the computer and the links. This is usually achieved by means of redundant data management on several mutually independent computers of a network. To 20 this end, the distribution of incoming messages to different servers by means of a so-called "load balancer" is known from the article "FGPA-based load balancer for Internet Servers, Harik, L., Kayssi, A." in Microelectronics, The 14th International Conference on 2002-ICM. The "load balancer" 25 thus forms a "single point of failure", i.e. when the "load balancer" itself fails, none of the subsequent servers can be reached.

Representation of the invention

30

The object underlying the invention is to specify an improved method for redundant data management in computer networks.

In accordance with the invention, this is carried out using a method for redundant data management in computer networks, with which data is consistently stored on at least two computers, with the application programs of the at least two 5 computers provided for data management comprising a corresponding address identifier, which features these application programs as a component of a virtual sub-network of a computer network and with the mutually independent computers having routing functions, by means of which the 10 application programs are accessed.

The invention can be used particularly advantageously in computer networks, in which communication is based on the Internet Protocol, and in which the Transmission Control 15 Protocol TCP or the User Datagram Protocol UDP is used as the transport protocol.

It is further favorable if an access ranking for the application programs provided for data management is 20 determined by the corresponding design of the routing function.

Brief description of the drawing

25 The invention is described in more detail with reference to a figure, which shows a schematic representation of addressing in an exemplary computer network.

Embodiment of the invention

30 Addressing in the illustrated computer network N is carried out according to the Internet Protocol version 4 by means of the so-called IP (Internet Protocol) address, which uniquely

identifies each network subscriber on level 3 of the OSI layer model, in other words on a logical level.

According to the Internet Protocol version 4, the IP address

5 is a 32 bit word comprising a first area, which is used to address the network and a second area, by means of which the computer in the network is identified.

The size of the two areas is standardized, a distinction being

10 made between 5 classes of IP addresses:

- Class A networks with 7 bit network addressing and 24 bits for the subscriber address
- Class B networks with 14 bit network addressing and 16 bit subscriber addressing
- Class C networks with 21 bit network addressing and 8 bit subscriber address
- Class D for addressing within the network with 28 bit subscriber addresses and
- Class E reserved for research purposes

20

The class of the IP address is shown in the leading bits.

Addresses of class A can be recognized here in that the leading bit comprises the value 0, class B by the value 10 of the two leading bits, class C by the 110 of the three first bits etc.

The routers of a network require the information about the address class in order to filter out the information about the address of the target network, with the aid of so-called

30 network masks, from the IP address of a target subscriber.

Whilst the network addresses are managed globally, the subscriber addresses are managed locally by the respective network operator. Said network operator is then also able to structure its network, by way of subdivision into sub-

networks, by corresponding allocation of subscriber addresses, without having to use any additional network numbers. To this end, the subscriber element of the IP address is subdivided into a

5. - sub-network address and the
- actual subscriber address.

The IP address is then made up of three parts, comprising the network address, the sub-network address and the actual
10 subscriber address.

A network of address class B can be used as a simple example, in which the IP address comprises the binary number 10 in the leading two bits and a further 14 bits for the address of the
15 network. The subscriber addressing element of 16 bits provided with this network class is in turn subdivided into an 8 bit element for addressing the sub-networks SN1, ... SN 25, ..., SN 255 and an 8 bit element for addressing the subscribers TN in the sub-networks. A maximum of 254 sub-networks can thus
20 be addressed with 254 subscribers respectively.

A typical IP address appears as follows in the generally conventional 4 octet notation of IP addresses, in which the individual bytes (8 bits) are shown by means of decimal
25 numerical values and are separated by periods:

172.16.115.54

The address of the network N itself here is 172.16.0.0. This
30 comprises up to 254 sub-networks SN1, ... SN 25, ..., SN 255 with the addresses

172.16.1.0, ..., 172.16.25.254.0

The 254 possible subscribers TN in the sub-network SN 25 with the address 172.16.25.0 then have the following addresses:

172.16.25.1, ..., 172.16.25.254

5

In accordance with the invention, the data to be managed in a redundant manner is consistently stored in two mutually independent computers R1, R2 of the network N, with data alignment between the redundant units being ensured using 10 known methods. These methods enable the two computers R1, R2 to belong to a network N of an operator or to different networks of the different operators connected to one another.

A corresponding address identifier is assigned to the

15 associated application programs of the two computers, said address identifier having these application programs as a component of a virtual sub-network SN 25 of a computer network N, such as the number 172.16.25.0. for instance. For the user's purposes, the application program operates on a 20 computer in this sub-network SN 25, for instance on a computer with the address 172.16.25.1.

This virtual sub-network SN 25 is connected logically to the network as a whole via two routers. To this end, the two

25 mutually independent computers R1, R2, which contain inter alia the application programs provided for the redundant and consistent storage of the data, are designed as routers, by means of which access to the application programs, the virtual sub-network, is exclusively available.

30

The remaining routers of the networks thus contain an indication in their routing tables that addresses of the type 172.16.25.XXX, which relate to the virtual sub-network SN 25, can only be reached via one of the two mentioned routers. In

this example, the entry in the routing table of the remaining routers contains the network address 172.16.25.0, the network mask 255.255.255.0, the costs for the path to this sub-network and the interface, by means of which this network N can be
5 reached.

All data packets with these target addresses thus reach one of the two routers and are forwarded from this to the corresponding application program with the respectively
10 associated storage area.

If network N is equipped with a dynamic routing method, as is generally conventional, the next query is automatically routed to the respectively favorable free storage area when one of
15 the two application program is busy. If the two routers, by means of which the virtual sub-network SN 25 can be reached, can be reached via equally expensive paths, a router is chosen at random and a load distribution is thus carried out.

20 In the case of a network with a static routing method, a corresponding load distribution must be carried out when determining the routing tables.

The specific exemplary embodiment describes the use of the
25 invention in a computer network N with addressing according to IP protocol version 4. The invention can however also be advantageously used for instance with IP protocol version 6. The use of the invention for all customary IP-based transport protocols, such as TCP or UDP, is likewise possible.

30 When using the TCP protocol, it must be ensured that access is carried out in each instance via a specific router. This can be achieved for instance via the cost functionality of dynamic routing functions.